arXiv:1509.08932v3 [cs.AI] 20 Oct 2015

Two Phase $Q-$learning for Bidding-based Vehicle Sharing

Yinlam Chow
Stanford University

Jia Yuan Yu
Concordia University

Marco Pavone
Stanford University

Abstract

We consider one-way vehicle sharing systems where customers can rent a car at one station and drop it off at another. The problem we address is to optimize the distribution of cars, and quality of service, by pricing rentals appropriately. We propose a bidding approach that is inspired from auctions and takes into account the significant uncertainty inherent in the problem data (e.g., pick-up and drop-off locations, time of requests, and duration of trips). Specifically, in contrast to current vehicle sharing systems, the operator does not set prices. Instead, customers submit bids and the operator decides whether to rent or not. The operator can even accept negative bids to motivate drivers to rebalance available cars to unpopular destinations within a city. We model the operator's sequential decision-making problem as a *constrained Markov decision problem* (CMDP) and propose and rigorously analyze a novel two phase Q-learning algorithm for its solution. Numerical experiments are presented and discussed.

1 Introduction

One-way vehicle sharing systems represent an increasingly popular mobility paradigm aimed at effectively utilizing usage of idle vehicles, reducing demands for parking spaces, and possibly cutting down excessive carbon footprints due to personal transportation. One-way vehicle sharing systems (also referred to as mobility-on-demand –MOD– systems) consist of a network of parking stations and a fleet of vehicles. A customer arriving at a given station can pick up a vehicle (if available) and drop it off at any other station within the city. Existing vehicle sharing systems include Zipcar [13], Car2Go [27] and Autoshare [25] for one-way car sharing, and Velib [21] and City-bike [8] for one-way bike sharing. Figure 1 shows a typical Toyota i-Road one-way vehicle sharing system [15].

Despite the apparent advantages, one-way vehicle sharing systems present significant operational challenges. Due to the asymmetry of travel patterns within a city, several stations will eventually experience imbalances of vehicle departures and customer arrivals. Stations with low customer demands (e.g., in suburbs) will possess excessive unused vehicles and require a large number of parking spaces, while stations with high demands (e.g., in the city center) will not be able to fulfill most customers' requests during rush hours.

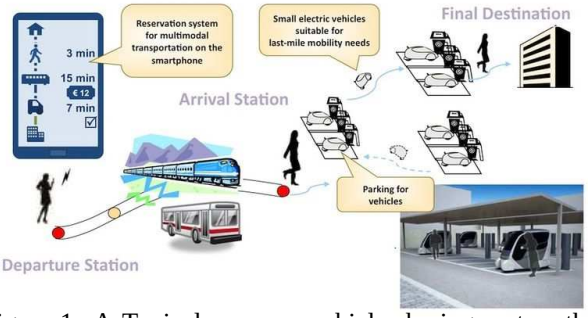

Figure 1: A Typical one-way vehicle sharing system that allows different pick-up and drop-off locations [15].

Literature Review: In general, there are two main methods available in the literature to address demand-supply imbalances in one-way vehicle sharing systems (however, if the vehicles can drive autonomously, additional rebalancing strategies are possible [34]). A first class of methods is to hire crew drivers to periodically relocate vehicles among stations. From a theoretical standpoint, optimal rebalancing of drivers and vehicles has been analyzed in [33], under the framework of queueing networks. In [2], [14], [28], the effectiveness of similar rebalancing strategies is numerically investigated via discrete event simulations. The work in [20] considers a stochastic mixed-integer programming model where the objective is to minimize vehicle relocation cost subject to a probabilistic constraint on the service level. The works in [30] and [18] consider a similar approach. While rebalancing can be quite effectively carried out by this way, these methods substantially increase sunk costs due to the large number of staff drivers that needs to be hired (a characterization of the number of drivers is provided in [33]), and may not scale well to large transportation networks.

Alternatively, demand-supply imbalances can be addressed by imposing incentive pricing to vehicle rentals. A typical incentive pricing mechanism is described in [23] and portrayed in Figure 2. The strategy is to adjust rental prices at each station as a function of current vehicle inventory and customers' requests. The work in [31] proposes a method to optimize vehicle assignments by trip splitting and trip

Preliminary work. Under review by AISTATS 2016. Do not distribute.

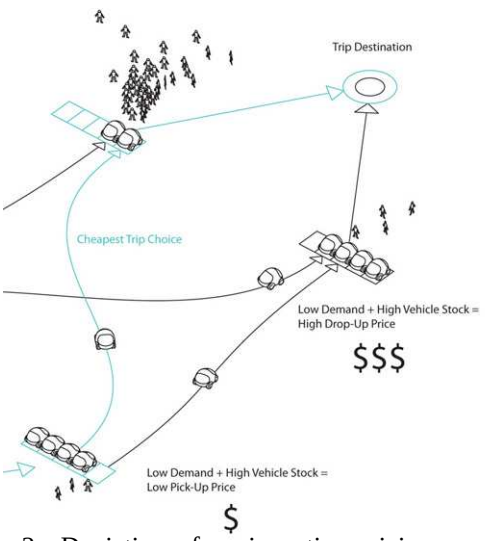

Figure 2: Depiction of an incentive pricing mechanism where rental prices are adjusted based on inventories and customers' demands [23].

joining, and [19] proposes a dynamic pricing strategy that enables clients of a carpooling system to trade-off convenience of a trip (i.e., duration) and cost. Carpooling strategies, however, may not be a scalable solution due to safety, convenience, and sociological reasons. Recently, the work in [6] proposes a bidding mechanism for vehicle rentals where at each station customers place bids and the vehicle sharing company decides which bids to accept. The operator's sequential decision-making problem is posed as a *constrained Markov decision problem* (CMDP), which is solved exactly, or approximately using an *actor-critic* method.

A bidding strategy, such as the one proposed in [6], is attractive for several reasons. First, accepted vehicle rental bids instantly reflect current demands and supplies at different stations. Second, by providing on-demand financial rewards for rebalancing vehicles, the rental company can save overhead costs associated with hiring crew drivers and renting extra parking spaces. Third, such pricing mechanism promotes high vehicle utilizations by encouraging extra vehicle rentals to less popular destinations and during non-rush hours.

Contributions: The contribution of this paper is threefold.

- Leveraging our recent findings in [6], we propose a novel CMDP formulation for the problem of optimizing one-way vehicle sharing systems. The actions in the CMDP model represent vehicle rental decisions as a function of bids placed by customers. The objective is to maximize total revenue from rental assignments, subject to a vehicle utilization constraint [1].
- We derive a two-phase Bellman optimality condition for CMDPs. Such condition shows that CMDPs can be solved exactly by two-phase dynamic programming (DP).
- We propose a novel sampling-based two-phase $Q-$learning algorithm for the solution of CMDPs, and show under mild assumptions convergence to an optimal solution.

This paper provides a first step toward designing market-based mechanisms for the operational optimization of one-way vehicle sharing systems. We describe a wealth of open problems at the end of the paper. Furthermore, the results concerning CMDPs are of independent interest and applicable more broadly. Due to space limitations, in this paper we only include the statements of our theoretical results. All proofs can be found in the supplementary material section.

2 Mathematical Model

In this section we present a mathematical model for one-way vehicle sharing systems and then pose a CMDP decision-making problem for its optimization.

2.1 General Model and Problem Data

Assume the vehicle sharing company owns C vehicles, indexed $1,\dots,C$, that can be parked at S stations, indexed $1,\dots,S$. The company only allows passenger to rent for a maximum of $\overline{T}$ time slots; furthermore, the maximum fare/reward for each rental period is $\overline{F}$. We consider a discrete time model $t = 0,1,2,\dots,T$, where T is the time horizon. At time $t \geq 0$ and at each station $j \in \{1,\dots,S\}$, customers' destinations, rental durations, and proposed travel fares can be modeled via a multi-variate (three-dimensional) stationary probability distributions Φ_t^j with domain $\{1,\dots,S\} \times (0,\overline{T}] \times [-\overline{F},\overline{F}]$. The distribution Φ_t^j models the inherent stochastic, time-variant customer demand for car rentals. A practical approach to estimate the demand distribution Φ_t^j at each station j is to use data-driven methods, for example [10].

At each time t and station $j \in \{1,\dots,S\}$, there are M_t^j rental requests, where M_t^j is modeled as a Poisson random variable with rate λ_t^j (as usually done in the literature, see, for example, [17]). The associated customers' destinations, rental durations, and proposed travel fares are i.i.d. samples drawn from the distribution Φ_t^j. Such samples are collected in a random vector ω_t^j, that is:

$$\omega_t^j = ((\mathbf{G}_t^{1,j},\mathbf{T}_t^{1,j},\mathbf{F}_t^{1,j}),\dots,(\mathbf{G}_t^{M_t^j,j},\mathbf{T}_t^{M_t^j,j},\mathbf{F}_t^{M_t^j,j})),$$

where random variables $\mathbf{G}$, $\mathbf{T}$, and $\mathbf{F}$ denote, respectively, a customer's destination, rental duration, and proposed fare (the meaning of the indices is clear from the definitions).

For $j \in \{1,\dots,S\}$, denote by $\mathcal{A}_t^{j,j'}$ the number of customers that arrive at time t at station j and that wish to travel to station j'. By definition, $\mathcal{A}_t^{j,j'} = \sum_{m=1}^{M_t^j} \mathbf{1}\{\mathbf{G}_t^{m,j} = j'\}$, where $1\{\cdot\}$ represents the indicator function.

[1] The vehicle utilization constraint ensures that the assignment policy does not excessively favor short-term rentals. Accordingly, this constraint provides service guarantees for customers in need of long-term rentals. Further details about the practical relevance of this constraint can be found in [5].

This model captures both concepts of renting and rebalancing. Note that at station j, the random price offered by customer m, i.e., $\mathbf{F}_t^{m,j}$, can either be positive or negative. When this quantity is positive, it means that the customer is willing to pay $\mathbf{F}_t^{m,j}$ fare units for a vehicle, traveling to station $\mathbf{G}_t^{m,j}$ in $\mathbf{T}_t^{m,j}$ time units. If this quantity is negative, it means that the company is paying $\mathbf{F}_t^{m,j}$ fare units to customer m to travel to station $\mathbf{G}_t^{m,j}$ in $\mathbf{T}_t^{m,j}$ time units to fulfill rebalancing needs.

Note that at each time instant and at each station, there can potentially be more rental requests than vehicles available. The strategy is to rank all incoming rental requests by a price-to-travel time function and assign vehicles according to such ranking. Specifically, for destination $j' \in \{1, \ldots, S\}$, we define the price-to-travel time function as

$$f_{\text{rank}}^{j'}(\mathbf{G}, \mathbf{T}, \mathbf{F}) = \begin{cases} \mathbf{1}\{\mathbf{F} \geq 0\}\mathbf{F}/\mathbf{T} + \mathbf{1}\{\mathbf{F} \leq 0\}\mathbf{F}\mathbf{T} & \text{if } \mathbf{G} = j' \\ -\infty & \text{otherwise} \end{cases}$$

By assigning vehicles according to the price-to-travel time function (in a descending order), one favors customers with high rental prices and short travel times (when renting), and drivers with low rewards and short travel times (when rebalancing). (If all vehicles move at similar speeds, one could equivalently consider distances instead of times.)

To explicitly model the instantaneous demand-supply imbalance throughout the transportation network, one can refine function $f_{\text{rank}}^{j'}(\mathbf{G}, \mathbf{T}, \mathbf{F})$ by including dependencies on the arrival station j and time instant t. While this generalization is straightforward, we omit the details in the interest of brevity.

2.2 State Variables

We now proceed to construct a CMDP model for a one-way vehicle sharing system. We consider the following state variables:

- For $t \geq 0$, $k_t = t \in \{0, \ldots, T-1\}$ is a counter state.
- For $i \in \{1, \ldots, C\}$ and $t \geq 0$, $q_t^i \in \{1, \ldots, S\}$ is the destination station at time t of the i^{th} vehicle. Let $q_t := (q_t^1, \ldots, q_t^C)$.
- For $i \in \{1, \ldots, C\}$ and $t \geq 0$, $\tau_t^i \in \{0, 1, 2, \ldots, \overline{T}\}$ is the remaining travel time to reach the destination for the i^{th} vehicle. Let $\tau_t := (\tau_t^1, \ldots, \tau_t^C)$.

Collectively, the state space is defined as $\mathbf{X} = \{0, \ldots, T-1\} \times \{1, \ldots, S\}^C \times \{0, 1, 2, \ldots, \overline{T}\}^C$. We let $x_0 = (0, q_0, \tau_0)$ denote the initial state, and $x_t = (k_t, q_t, \tau_t)$ the state at time t.

2.3 Decision Variables

At each time t, in order to maximize expected revenue, the company makes decisions about renting vehicles to customers. Specifically, at each time t, the decision variables are:

- For each station $j \in \{1, \ldots, S\}$, each vehicle $i \in \{1, \ldots, C\}$, and $t \geq 0$, $u_t^{i,j} \in \{0, 1\}$ is a binary decision variable that indicates whether vehicle i is destined to station j at time t. Let $u_t := (u_t^{1,1}, \ldots, u_t^{1,S}, \ldots, u_t^{C,1} \ldots, u_t^{C,S})$.

For each station-destination pair $(j, j') \in \{1, \ldots, S\}^2$, we consider the following constraint to upper bound the number of vehicle dispatches at time $t \geq 0$:

$$\begin{aligned} &\sum_{i=1}^{C} u_t^{i,j'} \mathbf{1}\{q_t^i = j,\, \tau_t^i = 0\} \leq \mathcal{A}_t^{j,j'} \text{ if } j \neq j', \\ &\sum_{i=1}^{C} u_t^{i,j} \mathbf{1}\{q_t^i = j,\, \tau_t^i = 0\} \leq \sum_{i=1}^{C} \mathbf{1}\{q_t^i = j,\, \tau_t^i = 0\}. \end{aligned} \tag{1}$$

Intuitively, constraints (1) restrict the number of vehicle dispatches to be less than the number of customer requests. In the special case when $j = j'$, the upper bound is $\sum_{i=1}^{C} \mathbf{1}\{q_t^i = j,\, \tau_t^i = 0\}$ instead of $\mathcal{A}_t^{j,j}$ because one needs to take into account the case when vehicle i stays idle at station j. Additionally, we consider the following constraints to guarantee well-posedness of the vehicles assignments:

$$\begin{aligned} &u_t^{i,j} = 1,\ \forall i \in \{1, \ldots, C\},\ \text{if } \tau_t^i > 0 \text{ and } q_t^i = j, \\ &\sum_{j=1}^{S} u_t^{i,j} = 1,\ \forall i \in \{1, \ldots, C\}. \end{aligned} \tag{2}$$

Accordingly, the action space is $\mathbf{U} = \{0, 1\}^{C \times S}$, and u_t is the action taken at time t. Furthermore, define the set of admissible controls at state $x \in \mathbf{X}$ as $\mathbf{U}(x) \subseteq \mathbf{U}$, such that $\mathbf{U}(x) = \{u \in \mathbf{U} \text{ and it satisfies constraints (1) and (2)}\}$.

2.4 State Dynamics

For each time t and station-destination pair $(j, j') \in \{1, \ldots, S\}^2$, let

$$\mathcal{S}_t(j, j') = \{i \in \{1, \ldots, C\} : \tau_t^i = 0, q_t^i = j, u_t^{i,j'} = 1\},$$

denote the set of vehicles at location j which are allocated to destination j'. For $j \neq j'$, denote the rental durations and proposed fares of the $|\mathcal{S}_t(j, j')|$ highest ranked customers (according to the price-to-travel time function $f_{\text{rank}}^{j'}(\cdot)$) as $\{(\mathbf{T}_{t,j,j'}^m, \mathbf{F}_{t,j,j'}^m)\}_{m=1}^{|\mathcal{S}_t(j,j')|}$.

The case $j = j'$ is special, as the quantity $|\mathcal{S}_t(j, j)|$ comprises vehicles that (1) will eventually return to station j, and (2) stay idle at station j. Notice that the definition of $u_t^{i,j}$ does not allow one to distinguish between these two cases. However, in this case, it is not profitable to rent vehicles to customers submitting negative bids. Hence one concludes that the number of vehicles that will eventually return to station j should be equal to the number of customers that want to return to station j and submit a positive bid. We denote such number as $V_t^{j,\text{in transit}}$. We similarly

define the vector $\{(\mathbf{T}^m_{t,j,j'}, \mathbf{F}^m_{t,j,j'})\}_{m=1}^{|\mathcal{S}_t(j,j')|}$, with the understanding that the first $V_t^{j,\text{in transit}}$ elements correspond to the customers that submit positive bids, while each remaining element is set to $(0,0)$.

The state dynamics are then given as follows:

- For each vehicle $i \in \{1,\dots,C\}$ such that $\tau_t^i > 0$,
$$(k_{t+1}, q_{t+1}^i, \tau_{t+1}^i) = (k_t + 1, q_t^i, \max(\tau_t^i - 1, 0)).$$
- For all other vehicles, for each station-destination pair $(j,j') \in \{1,\dots,S\}^2$, and each vehicle $i \in \mathcal{S}_t(j,j')$, uniformly randomly allocate $\{\mathbf{T}^1_{t,j,j'},\dots,\mathbf{T}^{|\mathcal{S}_t(j,j')|}_{t,j,j'}\}$ to the available vehicles, i.e.,
$$(k_{t+1}, q_{t+1}^i, \tau_{t+1}^i) = \left(k_t + 1, j', \mathbf{T}^m_{t,j,j'}\right),$$
for all $m \in \{1,\dots,|\mathcal{S}_t(j,j')|\}$.

As the problem is high-dimensional even for a moderate number of vehicles and stations, and since the state update equations depend in a rather involved fashion on the information vector ω_t^j, the explicit derivation of the state transition probabilities (denote by $\mathbb{P}$) is impractical. This motivates the Q-learning approach proposed in this paper.

2.5 Revenue and Constraint Cost Functions

Define the *immediate* reward function $\mathbf{R} : \mathbf{X} \times \mathbf{U} \to \mathbb{R}$ as

$$\mathbf{R}(x,u) = \mathbb{E}\left[\sum_{j=1}^{S}\sum_{j'=1}^{S}\sum_{m=1}^{|\mathcal{S}_t(j,j')|}\mathbf{F}^m_{t,j,j'}\right].$$

The total revenue is then given by

$$\sum_{t=0}^{T}\mathbb{E}\left[\mathbf{R}(x_t,u_t)\right].$$

From a pure profit maximization standpoint, an operator should favor short term rental assignments, in order to minimize the opportunity cost of rejecting future customers that might potentially be more profitable. While this strategy optimizes the long-term revenue, it is not fair toward customers that require extended rental periods. To balance total profit with customers' satisfaction, we impose an additional constraint that lower bounds vehicle utilization. Specifically, the constraint cost function $\mathbf{D} : \mathbf{X} \times \mathbf{U} \to \mathbb{R}$ captures vehicle utilization at each time instant according to the definition

$$\mathbf{D}(x,u) = \mathbf{d} - \sum_{i=1}^{C}\frac{\tau^i}{TC},$$

where $\mathbf{d}$ represents a threshold for the average utilization rate specified by the system operator. The vehicle utilization constraint is then given by

$$\sum_{t=0}^{T}\mathbb{E}\left[\mathbf{D}(x_t,u_t)\right] \le 0.$$

The objective is then to maximize expected revenue while satisfying the vehicle utilization constraint.

2.6 CMDP Formulation

Equipped with the state space $\mathbf{X}$, control space $\mathbf{U}$, immediate reward $\mathbf{R}$, transition probability $\mathbb{P}$ (implicitly defined above), initial state x_0, and immediate constraint cost function $\mathbf{D}$, we pose the problem of controlling a vehicle sharing system as a CMDP. Specifically, note that the problem can be formulated on an infinite horizon by interpreting the set of states $\mathcal{X} = \{x_t \in \mathbf{X} : k_t = T\}$ (indeed containing a single state) as *absorbing*, i.e., once $k_t = T$, the system enters a termination state, call it `END`, and stays there with *zero* reward and constraint cost. From this perspective, T represents a *deterministic* stopping time, and the upper limits in the summations for the reward and cost functions can be replaced with ∞. We also define $\mathbf{X}' = \mathbf{X} \setminus \mathcal{X}$ as the set of *transient states*. Let Π_M be the set of closed-loop, Markovian stationary policies $\mu : \mathbf{X} \to \mathbb{P}(\mathbf{U})$. A policy $\mu \in \Pi_M$ induces a stationary mass distribution over the realizations of the stochastic process (x_t, u_t). It is well known that for CMDPs there is no loss of optimality in restricting the attention to policies in Π_M (instead, e.g., of also considering history-dependent policies).

Accordingly, in this paper we wish to solve the CMDP problem:

Problem $\mathcal{OPT}$ – Solve

$$\begin{aligned}\text{maximize}_{\mu\in\Pi_M} \quad & \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{R}(x_t,u_t) \mid x_0\ u_t \sim \mu\right] \\ \text{subject to} \quad & \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t) \mid x_0\ u_t \sim \mu\right] \le 0.\end{aligned}$$

Since, in our problem, the state and action spaces are exponentially large (in terms of the vehicle number C and station number S) and the explicit derivation of the state transition probability is intractable, exact solution methods for CMDPs (e.g., [1]) are not applicable.

On the other hand, even if one computes an optimal randomized policy for Problem $\mathcal{OPT}$, executing such policy on a vehicle sharing platform with simple system architecture may lead to problems in vehicle mis-coordination [24]. This motivates us to restrict the structure of admissible policies in problem $\mathcal{OPT}$ to Markovian, stationary and deterministic. In the next section we will introduce a two-phase Bellman optimality condition for problem $\mathcal{OPT}$, which serves as the theoretical underpinning for the design of an asymptotically optimal two-phase $Q-$learning algorithm.

3 Two-Phase Dynamic Programming Algorithm

In this section, by leveraging the results in [11], we present a two-phase DP algorithm for problem $\mathcal{OPT}$. As we shall see, the first step is to compute a Q-function that allows one to *refine* the set of feasible control actions (essentially, we retain only those actions that can guarantee the fulfillment of the constraint in problem $\mathcal{OPT}$). We then define a Bellman operator (restricted to the refined set of control actions) that allows the computation of an optimal policy. The DP algorithm presented in this section provides the

conceptual basis for the two-phase $Q-$learning algorithm presented in the next section.

3.1 Phase 1: Finding the Feasible Set

In this section we characterize the feasible set for problem $\mathcal{OPT}$ using the set of optimal policies from an auxiliary MDP problem, defined as follows:

Problem $\mathcal{FS}$ – Solve

$$\min_{\mu\in\Pi_{MD}} \max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\right\}, \tag{3}$$

where Π_{MD} is the set of closed-loop, Markovian, stationary, and deterministic policies.

The following (trivial) result related minimizer for problem $\mathcal{FS}$ with feasible solutions for problem $\mathcal{OPT}$.

Lemma 1 *Let $\mu : \mathbf{X} \to \mathbf{U}$ be a Markovian stationary deterministic policy that minimizes problem $\mathcal{FS}$, that is*

$$\mu \in \arg\min_{\mu} \max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\right\}.$$

Then the solution cost is equal to zero if and only if

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right] \le 0.$$

Equipped with the above result, one can immediately deduce the following:

- If the solution cost to problem $\mathcal{FS}$ is strictly larger than zero, i.e.,

$$\min_{\mu\in\Pi_{MD}} \max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\right\} > 0,$$

  then problem $\mathcal{OPT}$ is infeasible.

- Otherwise, the feasible set of Markovian stationary deterministic policies is given by

$$\Pi_{\text{FS}} = \arg\min_{\mu\in\Pi_{MD}} \max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\right\}.$$

In order to characterize the feasible set Π_{FS}, we derive a Bellman optimality condition for problem $\mathcal{FS}$ and demonstrate how Π_{FS} can be computed via DP.

Before getting into the main result, we define the Bellman operator for problem $\mathcal{FS}$ as follows:

$$\mathbf{T}[V](x) := \min_{u\in\mathbf{U}(x)} \max\left\{\mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V(x')\right\}$$

where $\mathcal{B}(x)$ is the indicator function

$$\mathcal{B}(x) = \begin{cases} 0 & \text{if } x\in\mathcal{X} \\ -\infty & \text{otherwise} \end{cases}.$$

With such definition of Bellman operator $\mathbf{T}$, we will later see that the fixed point solution of $\mathbf{T}[V](x) = V(x), \forall x \in \mathbf{X}'$ is equal to the solution of problem $\mathcal{FS}$, given by (3).

For any bounded initial value function estimate $V_0 : \mathbf{X} \to \mathbb{R}$ with $V_0(x) = 0$ for $x \in \mathcal{X}$, we define the value function sequence

$$V_{k+1}(x) = \mathbf{T}[V_k](x),\ \forall x\in\mathbf{X}',\ k\in\{0,1,\ldots\}. \tag{4}$$

The following theorem shows that this sequence of value function estimates converges to the solution of problem $\mathcal{FS}$, which is also the unique fixed point of $\mathbf{T}[V](x) = V(x), \forall x \in \mathbf{X}'$.

Theorem 2 (Bellman Optimality for $\mathcal{FS}$) *For any initial value function estimate $V_0 : \mathbf{X} \to \mathbb{R}$ where $V_0(x) = 0$ at $x \in \mathcal{X}$, there exists a limit function V^* such that*

$$\begin{aligned} V^*(x_0) &= \lim_{N\to\infty} \mathbf{T}^N[V_0](x_0) \\ &= \min_{\mu\in\Pi_M} \max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\,|\,x_0,\mu\right]\right\}. \end{aligned} \tag{5}$$

Furthermore, V^* is a unique solution to the fixed point equation: $\mathbf{T}[V](x) = V(x), \forall x \in \mathbf{X}'$.

By running the value iteration algorithm in (4), one obtains the optimal value function for problem $\mathcal{FS}$. If $V^*(x_0) = 0$, then, by Lemma 1, every feasible policy for problem $\mathcal{OPT}$, denoted by μ_{FS}, can be obtained as

$$\mu_{\text{FS}}(x) \in \arg\min_{u\in\mathbf{U}(x)} \max\left\{\mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V^*(x')\right\}.$$

However, since the number of feasible policies is exponential in the size of state and action spaces, their exhaustive enumeration is intractable, and the above result is useful only form a conceptual standpoint. To address this problem, we consider a refined notion of feasible control actions in terms of optimal state-action value functions ($Q-$functions), which provides the basis for the two-phase DP algorithm. Specifically, a $Q-$function is defined as:

$$Q^*(x,u) := \max\left\{\mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V^*(x')\right\},$$

where $V^*(x) = \min_{u\in\mathbf{U}(x)} Q^*(x,u)$ for $x \in \mathbf{X}'$. By defining the state-action Bellman operator

$$\mathbf{F}[Q](x,u) = \max\left\{\mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)\min_{u'\in\mathbf{U}(x')} Q(x',u')\right\},$$

equivalently Q^* is a unique fixed point solution to $\mathbf{F}[Q](x,u) = Q(x,u)$ for any $u \in \mathbf{U}(x)$, $x \in \mathbf{X}'$. Note the interpretation of $Q^*(x,u)$ from these equations: it is the constraint cost of starting

at state $x \in \mathbf{X}'$, using control action u in the first stage, and using an optimal policy of problem $\mathcal{FS}$ thereafter.

We are now in a position to define a *refined* set of feasible control actions, denoted by $\mathbf{U}_{\text{FS}}(Q^*, x)$, whereby we retain only those actions that can guarantee the fulfillment of the constraint in problem $\mathcal{OPT}$. Specifically, for any state $x \in \mathbf{X}'$, we define:

$$\mathbf{U}_{\text{FS}}(Q^*, x) := \{u \in \mathbf{U}(x) : \exists \mu_{\text{FS}} \in \Pi_{\text{FS}} \text{ such that } u = \mu_{\text{FS}}(x)\}$$
$$= \left\{ u \in \mathbf{U}(x) : u \in \arg \min_{u' \in \mathbf{U}(x)} Q^*(x, u') \text{ and } Q^*(x, u) = 0 \right\}.$$

3.2 Phase 2: Constrained Optimization

From the definition of the set of feasible policies Π_{FS}, one can reformulate problem $\mathcal{OPT}$ as:

$$\max_{\mu \in \Pi_{\text{FS}}} \mathbb{E}\left[\sum_{t=0}^{T} \mathbf{R}(x_t, u_t) \mid x_0, \mu\right]. \tag{6}$$

This problem can be solved via value iteration by defining the Bellman operator (with respect to the *refined* set of control actions $\mathbf{U}_{\text{FS}}(Q^*, x)$):

$$\mathbf{T}_R[W](x) := \max_{u \in \mathbf{U}_{\text{FS}}(Q^*, x)} \left\{ \mathbf{R}(x, u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x, u) W\left(x'\right) \right\}.$$

The following theorem shows there exists a unique fixed point solution to $\mathbf{T}_R[W](x) = W(x)$ and such solution corresponds to the value function for the problem in (6) (and, hence, problem $\mathcal{OPT}$).

Theorem 3 (Bellman Optimality for $\mathcal{OPT}$) *For any initial value function estimate $W_0 : \mathbf{X} \to \mathbb{R}$ such that $W_0(x) = 0$ for any $x \in \mathcal{X}$, there exists a limit function W^* such that*

$$W^*(x_0) = \lim_{N \to \infty} (\mathbf{T}_R)^N [W_0](x_0)$$
$$= \max_{\mu \in \Pi_{\text{FS}}} \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{R}(x_t, u_t) \mid x_0, \mu\right].$$

Furthermore, W^* is a unique solution to the fixed point equation: $\mathbf{T}_R[W](x) = W(x)$, for any $x \in \mathbf{X}'$.

Therefore, for any bounded initial value function estimate $W_0 : \mathbf{X} \to \mathbb{R}$ such that $W_0(x) = 0$ at $x \in \mathcal{X}$, the value function estimate sequence

$$W_{k+1}(x) = \mathbf{T}_R[W_k](x), \ \forall x \in \mathbf{X}', \ k \in \{0, 1, \ldots, \}, \tag{7}$$

converges to the value function for problem $\mathcal{OPT}$.

Finally, we define the $Q-$function for problem (6) (denoted as H to distinguish it from the $Q-$function for problem $\mathcal{FS}$):

$$H^*(x, u) := \mathbf{R}(x, u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x, u)\, W^*(x'),$$

where $W^*(x)$ is the value function for problem (6). By defining the state-action Bellman operator

$$\mathbf{F}_R[H](x, u) = \mathbf{R}(x, u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x, u) \min_{u' \in \mathbf{U}_{\text{FS}}(Q^*, x')} H\left(x', u'\right),$$

one can show, similarly as before, that H^* is a unique fixed point solution of $\mathbf{F}_R[H](x, u) = H(x, u)$ for $u \in \mathbf{U}_{\text{FS}}(Q^*, x)$, $x \in \mathbf{X}'$.

The above two-phase Bellman optimality condition immediately leads to a two-phase DP algorithm for the solution of problem $\mathcal{OPT}$. However, such DP algorithm presents two main implementation challenges. First, the algorithm is not applicable to the vehicle sharing problem considered in this paper since the state transition probabilities are not available explicitly (as discussed above). Second, when the size of the state and action spaces are large, updating the value iteration estimates is computationally intractable.

To address these computation challenges, in the next section we present a sampling-based two-phase $Q-$learning algorithm that approximates the solution to problem $\mathcal{OPT}$. Similar to the two-phase DP algorithm, in the first phase one updates the $Q-$function estimates for problem $\mathcal{FS}$ "by sampling" the vehicle sharing model. Then, in a second phase, one updates the $Q-$function estimates for problem $\mathcal{OPT}$ (recall that such $Q-$functions are referred to as functions H).

4 Two Phase $Q-$learning

In this section we present both synchronous and asynchronous versions of two-phase $Q-$learning to solve problem $\mathcal{OPT}$. In the synchronous version, the $Q-$function estimates of all state-action pairs are updated at each step. In contrast, in the asynchronous version, only the $Q-$function estimate of a sampled state-action pair is updated. Under mild assumptions, we show that both algorithms are asymptotically optimal. While the convergence rate of synchronous $Q-$learning is higher [12], asynchronous $Q-$learning is more computationally efficient.

4.1 Synchronous Two Phase $Q-$learning

Suppose $Q_0(x, u)$ is an initial $Q-$function estimate such that $Q_0(x, u) = 0$ for any $x \in \mathcal{X}$. At iteration $k \in \{0, 1, \ldots\}$, the synchronous two-phase $Q-$learning algorithm samples N states $(x'^{,1}, \ldots, x'^{,N})$ and updates the $Q-$function estimates for each state-action pair $(x, u) \in \mathbf{X}' \times \mathbf{U}$ as follows:

$$Q_{k+1}(x, u) = Q_k(x, u) + \zeta_{2,k}(x, u) \cdot \Bigg(\max \Bigg\{ \mathcal{B}(x), \mathbf{D}(x, u) + \frac{1}{N} \sum_{m=1}^{N} \min_{u'^{,m} \in \mathbf{U}(x'^{,m})} Q_k(x'^{,m}, u'^{,m}) \Bigg\} - Q_k(x, u) \Bigg), \tag{8}$$

$$H_{k+1}(x, u) = H_k(x, u) + \zeta_{1,k}(x, u) \cdot \Bigg(\mathbf{R}(x, u) + \frac{1}{N} \sum_{m=1}^{N} \max_{u'^{,m} \in \mathbf{U}_{\text{FS}}(Q_k, x'^{,m})} H_k(x'^{,m}, u'^{,m}) - H_k(x, u) \Bigg). \tag{9}$$

The step size pair $(\zeta_{1,k}(x, u), \zeta_{2,k}(x, u))$ follows the update rule

$$\sum_k \zeta_{1,k}(x, u) = \sum_k \zeta_{2,k}(x, u) = \infty,$$
$$\sum_k \zeta_{1,k}^2(x, u) < \infty, \ \sum_k \zeta_{2,k}^2(x, u) < \infty, \tag{10}$$
$$\zeta_{1,k}(x, u) = o\big(\zeta_{2,k}(x, u)\big).$$

The last equation implies that the $Q-$function update $Q_k(x, u)$ is on the fast time scale, and the $Q-$function update $H_k(x, u)$ is on the slow time scale. Notice that in the sampling approach, the state trajectory will enter the absorbing set $\mathcal{X}'$ in T steps. While the convergence of $Q-$learning is a polynomial of $|\mathbf{X}|$ and $|\mathbf{U}|$

(see the finite sample analysis in Theorem 1 of [12]), in order to get an accurate estimate of the $Q-$function one needs more state-action samples from the transient state space. However once the state trajectory enters $\mathcal{X}$, it will never visit the transient state space $\mathbf{X}'$ again. To collect more samples from the transient state space, similar to the approaches adopted by sampling based methods in [32], [9], here we reset the state to its initial condition immediately after it enters the absorbing set. The convergence result for the synchronous two-phase $Q-$learning algorithm is given in the following theorem.

Theorem 4 (Convergence of Synchronous Q-learning) *Suppose the step-sizes $(\zeta_{1,k}(x,u), \zeta_{2,k}(x,u))$ follow the update rule in* (10)*. Then the sequence of $Q-$function estimates computed via synchronous two-phase $Q-$learning converges to the optimal $Q-$function pair $(Q^*(x,u), H^*(x,u))$ component-wise with probability* 1.

After both $Q-$functions converge, a near-optimal policy can be computed as

$$\widetilde{\mu}^*(x) \in \arg \min_{u \in \mathbf{U}_{\mathrm{FS}}(Q_{\bar{k}},x)} H_{\bar{k}}(x,u), \ \forall x \in \mathbf{X}', \qquad (11)$$

where $\bar{k}$ is the iteration index when the leaning is stopped.

4.2 Asynchronous Two-Phase $Q-$learning

Suppose $Q_0(x,u)$ is an initial $Q-$function estimate such that $Q_0(x,u) = 0$ for any $x \in \mathcal{X}$. At iteration $k \in \{0,1,\ldots\}$ and state $x_k \in \mathbf{X}$, the asynchronous two-phase $Q-$learning algorithm (1) generates a control action

$$u_k \in \arg \min_{u \in \mathbf{U}_{\mathrm{FS}}(Q_k,x_k)} H_k(x_k,u), \qquad (12)$$

(2) samples N states $(x'^{,1}, \ldots, x'^{,N})$, and (3) updates the $Q-$function estimates as follows:

- for $x = x_k$ and $u = u_k$, $Q-$function estimates are updated according to equations (8) and (9),
- otherwise, the $Q-$function estimates are equal to their previous values, i.e.,

$$Q_{k+1}(x,u) = Q_k(x,u), \ H_{k+1}(x,u) = H_k(x,u).$$

The convergence result for the asynchronous two-phase $Q-$learning algorithm is given in the following theorem.

Theorem 5 (Convergence of Asynchronous Q-learning) *Suppose the step-sizes $(\zeta_{1,k}(x,u), \zeta_{2,k}(x,u))$ follow the update rule in* (10)*. Also suppose each state action pair $(x,u) \in \mathbf{X} \times \mathbf{U}$ is visited infinitely often. Then, the sequence of $Q-$function estimates computed via asynchronous two-phase $Q-$learning converges to the optimal $Q-$function pair $(Q^*(x,u), H^*(x,u))$ with probability* 1.

Note that the convergence result relies on the assumption that each state-action pair $(x,u) \in \mathbf{X} \times \mathbf{U}$ is visited infinitely often. While this is a standard assumption in the $Q-$learning literature [3], by following analogous arguments as in [29], the above result can proven under milder assumptions by using PAC analysis. As for synchronous two-phase $Q-$learning, a near optimal policy can be computed by (11) after the $Q-$functions converge.

In the next section, we perform numerical experiments to compare the proposed two-phase $Q-$learning algorithms to a number of alternative approaches. In particular, we consider a Lagrangian relaxation method [1], whereby one transforms problem $\mathcal{OPT}$ into a min-max MDP and solve for the optimal saddle point. However, finding the optimal Lagrange multiplier is a challenging problem. While multi-scale stochastic approximation algorithms such as actor-critic [4] are available for optimizing both the Lagrange multiplier and policy online, in order to update the Lagrange multiplier, one requires sequential gradient approximations. This makes the convergence of these algorithms very sensitive to the multiple step-sizes and thus non-robust to large scale problems. Further numerical insights are provided below.

4.3 Numerical Results

Consider a small vehicle sharing system that consists of 15 vehicles ($C = 15$), 5 stations ($S = 5$), and a horizon of 6 hours ($T = 6$). In problem $\mathcal{OPT}$, one aims to find an optimal assignment strategy that maximizes the total revenue subject to the vehicle utilization constraint. The constraint threshold $\mathbf{d}$ is set equal to $38 \times T$ to ensure that the average utilization time of each vehicle is at least 2.5 hours. In our comparative study, we consider (1) a $Q-$learning algorithm [12] that maximizes total revenue and does not take into account the utilization constraint; (2) a penalized $Q-$learning algorithm, which maximizes a combined utility function of revenue and constraint violation penalty; (3) a $Q-$learning algorithm with Lagrangian update [4], which approximately solves problem $\mathcal{OPT}$ using an actor-critic method; and (4) the proposed two-phase $Q-$learning algorithm. Performance of these algorithms is evaluated via 1000 Monte Carlo trials, for which the corresponding empirical rewards and constraint costs are shown in Figures 3 and 4 respectively. The policy computed via $Q-$learning returns the highest total revenue, but the average vehicle utilization time is only 1.7 hours. On the other hand, the computed policies from two-phase $Q-$learning and $Q-$learning with Lagrangian update decrease total revenue by 20% but guarantee that average vehicle utilization time is over 2.5 hours. Also, one can note that the proposed two-phase $Q-$learning algorithm converges faster than the $Q-$learning algorithm with Lagrangian update. Finally, the policy from penalized $Q-$learning[2] has the highest average vehicle utilization time (2.73 hours) but lowest total revenue (with a 28% gap). As a further comparison, a greedy policy that assigns as many rentals as possible provides a total revenue as low as 33.05.

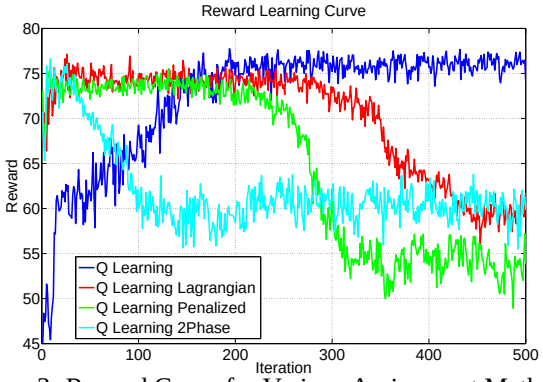

Figure 3: Reward Curve for Various Assignment Methods.

In each iteration of the two-phase $Q-$learning algorithm, the assignment problem in (12) is cast as a *bilinear integer linear programming* (BILP) problem. Although BILP problems are NP-hard in general, with readily available optimization packages such as CPLEX [7], our algorithm is capable of solving medium-scale problems with up to 50 vehicles, 20 stations, and a horizon of 12 hours. We believe there is still ample room for improvement by leveraging parallelization and characterizing the $Q-$functions with function approximations.

5 Conclusion

In this paper, we propose a novel CMDP model for one-way vehicle sharing systems whereby the real-time rental assign-

[2]Here we perform grid search on the penalty parameter in order to maximize the total revenue while satisfying the vehicle utilization constraint.

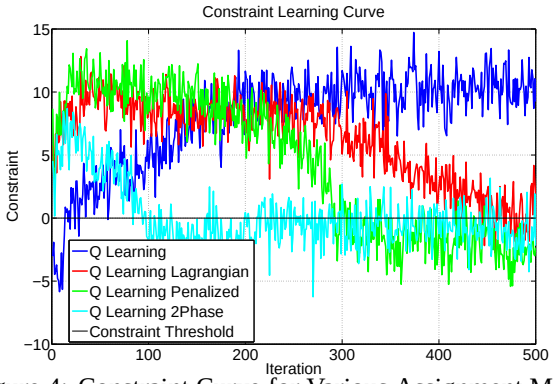

Figure 4: Constraint Curve for Various Assignment Methods.

ment of vehicles relies on an auction-style bidding paradigm. We rigorously derive a two-phase Bellman optimality condition for the CMDP and show that this problem can be solved (conceptually) using two-phase dynamic programming. Building upon this result, we propose a practical sampling-based two-phase $Q-$learning algorithm and show that the solution converges asymptotically to the value function of the CMDP.

Future work includes: **1)** Providing a convergence rate for our two-phase $Q-$learning algorithm; **2)** Generalizing the proposed bidding mechanism by using market-design mechanisms [16] and game theory [22]; and **3)** Evaluating our algorithm on a large-scale vehicle sharing system.

6 Acknowledgements

Y.-L. Chow and M. Pavone are partially supported by Croucher Foundation Doctoral Scholarship, the Office of Naval Research, Science of Autonomy Program, under Contract N00014-15-1-2673, and by the National Science Foundation under CAREER Award CMMI-1454737.

References

[1] E. Altman. *Constrained Markov Decision Processes*, volume 7. CRC Press, 1999.

[2] M. Barth and M. Todd. Simulation Model Performance Analysis of a Multiple Station Shared Vehicle System. *Transportation Research Part C: Emerging Technologies*, 7(4):237–259, 1999.

[3] D. Bertsekas and J. Tsitsiklis. *Neuro-dynamic Programming*. Athena Scientific, 1996.

[4] V. Borkar. An Actor-critic Algorithm for Constrained Markov Decision Processes. *Systems & Control letters*, 54(3):207–213, 2005.

[5] D. Brook. Carsharing and Vehicle Sharing Report. `http://fresno.ts.odu.edu/newitsd/ITS_Serv_Tech/carsharing/carsharing_report.html`, 2005. California Center for Innovative Transportation at the University of California at Berkeley and Caltrans.

[6] Y. Chow and J. Yu. Real-time Bidding Based Vehicle Sharing. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1829–1830. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[7] CPLEX, IBM ILOG. V12. 1: User's Manual for CPLEX. *International Business Machines Corporation*, 46(53):157, 2009.

[8] M. DiDonato. *City-bike Maintenance and Availability*. PhD thesis, Worcester Polytechnic Institute, 2002.

[9] D. Egloff et al. Monte Carlo Algorithms for Optimal Stopping and Statistical Learning. *The Annals of Applied Probability*, 15(2):1396–1432, 2005.

[10] V. Epanechnikov. Non-parametric Estimation of a Multivariate Probability Density. *Theory of Probability & Its Applications*, 14(1):153–158, 1969.

[11] Z. Gábor, Z. Kalmár, and C. Szepesvári. Multi-criteria Reinforcement Learning. In *ICML*, volume 98, pages 197–205, 1998.

[12] M. Kearns and S. Singh. Finite-sample Convergence Rates for $Q-$learning and Indirect Algorithms. *NIPS*, pages 996–1002, 1999.

[13] P. Keegan. Zipcar-The Best New Idea in Business. *CNN-Money. com*, 2009.

[14] A. Kek et al. Relocation Simulation Model for Multiple-station Shared-use Vehicle Systems. *Transportation Research Record: Journal of the Transportation Research Board*, 1986(1):81–88, 2006.

[15] J. Kendall. Toyota Takes the i-Road. *Automotive Engineering International*, 121(3), 2013.

[16] V. Krishna. *Auction Theory*. Academic press, 2009.

[17] Y. Lim, S. Lee, and J. Kim. A Stochastic Process Model for Daily Travel Patterns and Traffic Information. In *Agent and Multi-Agent Systems: Technologies and Applications*, pages 102–110. Springer, 2007.

[18] D. Mauro et al. The Bike Sharing Rebalancing Problem: Mathematical Formulations and Benchmark Instances. *Omega*, 45:7–19, 2014.

[19] W. Mitchell. *Reinventing the Automobile: Personal Urban Mobility for the 21st Century*. MIT press, 2010.

[20] R. Nair. Fleet Management for Vehicle Sharing Operations. *Transportation Science*, 45(4):524–540, 2011.

[21] R. Nair et al. Large-Scale Vehicle Sharing Systems: Analysis of Vélib. *International Journal of Sustainable Transportation*, 7(1):85–106, 2013.

[22] N. Nisan et al. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[23] D. Papanikolaou et al. *The Market Economy of Trips*. PhD thesis, Massachusetts Institute of Technology, 2011.

[24] P. Paruchuri, M. Tambe, F. Ordonez, and S. Kraus. Towards a Formalization of Teamwork with Resource Constraints. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 596–603. IEEE Computer Society, 2004.

[25] E. Reynolds and K. McLaughlin. Autoshare: The Smart Alternative to Owning a Car. *Autoshare, Toronto, Ontario, Canada*, 2001.

[26] S. Ross. *Stochastic Processes*, volume 2. John Wiley & Sons New York, 1996.

[27] A. Schmauss. Car2go in Ulm, Germany, as an Advanced Form of Car-sharing. *European Local Transport Information Service (ELTIS)*, 2009.

[28] J. Shu et al. Bicycle-sharing System: Deployment, Utilization and the Value of Re-distribution. *National University of Singapore-NUS Business School, Singapore*, 2010.

[29] A. Strehl, L. Li, E. Wiewiora, J. Langford, and M. Littman. PAC Model-free Reinforcement Learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 881–888. ACM, 2006.

[30] R. Tal et al. Static Repositioning in a Bike-sharing System: Models and Solution Approaches. *European Journal of Transportation and Logistics*, 2:187–229, 2013.

[31] K. Uesugi et al. Optimization of Vehicle Assignment for Car Sharing System. In *Knowledge-based intelligent information and engineering systems*, pages 1105–1111. Springer, 2007.

[32] H. Yu and D. Bertsekas. A Least Squares $Q-$learning Algorithm for Optimal Stopping Problems. *Lab. for Information and Decision Systems Report*, 2731, 2006.

[33] R. Zhang and M. Pavone. A Queueing Network Approach to the Analysis and Control of Mobility-on-demand Systems. In *Proceedings of American Control Conference*, 2015.

[34] R. Zhang and M. Pavone. Control of Robotic Mobility-On-Demand Systems: A Queueing-Theoretical Perspective. *International Journal of Robotics Research*, 2015.

A Appendix: Technical Proofs

A.1 Proof of Lemma 1

First notice that

$$\max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\right\}\geq 0.$$

Thus for any minimizer μ^* of problem $\mathcal{FEA}$ such that the solution is 0, it directly implies that

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu^*\right]\leq 0,$$

i.e., μ^* is a feasible policy of problem $\mathcal{OPT}$.

On the other hand, suppose a control policy μ is feasible to problem $\mathcal{OPT}$, i.e.,

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\leq 0.$$

This implies that

$$\max\left\{0, \mathbb{E}\left[\sum_{t=0}^{\infty}\mathbf{D}(x_t,u_t)\mid x_0,\mu\right]\right\}= 0,$$

Therefore μ is a minimizer to problem $\mathcal{FEA}$ because the objective function of this problem is always non-negative.

A.2 Technical Properties of Bellman Operators

The Bellman operator $\mathbf{T}$ has the following properties.

Lemma 6 *The Bellman operator $\mathbf{T}[V]$ has the following properties:*

- *(Monotonicity) If $V_1(x)\geq V_2(x)$, for any $x\in\mathbf{X}$, then $\mathbf{T}[V_1](x)\geq\mathbf{T}[V_2](x)$.*
- *(Translational Invariant) For any constant $K\in\mathbb{R}$, $\mathbf{T}[V](x)-|K|\leq\mathbf{T}[V+K](x)\leq\mathbf{T}[V](x)+|K|$, for any $x\in\mathbf{X}'$.*
- *(Contraction) There exists a positive vector $\{\xi(x)\}_{x\in\mathbf{X}}$ and a constant $\beta\in(0,1)$ such that $\|\mathbf{T}[V_1]-\mathbf{T}[V_2]\|_\xi\leq\beta\|V_1-V_2\|_\xi$*[3]*.*

Proof 1 *The proof of monotonicity and constant shift properties follow directly from the definition of Bellman operator. Now we prove the contraction property. Recall that the $t-$element in state $x=(t,z,\omega)$ is a time counter, its transition probability is given by $\mathbf{1}\{t'=t+1\}$ if $t<T-1$ and $\mathbf{1}\{t'=t\}$ if $t=T-1$. Obviously the transition probability $\mathbb{P}(x'|x,u)$, which is a multivariate probability distribution of state x, is less than or equal to the marginal probability distribution of $t-$element. Thus for vector $\{\xi(x)\}_{x\in\mathbf{X}}$ such that*

$$\xi(x)=T-t\geq 0,\ \forall x\in\mathbf{X}, \tag{13}$$

we have that

$$\sum_{x'\in\mathbf{X}'}\xi(x')\mathbb{P}(x'|x,u)\leq\sum_{x'\in\mathbf{X}'}\xi(x')\mathbf{1}\{t'=t+1\}\leq\frac{T-1}{T}\xi(x),\ \forall x\in\mathbf{X},\ \forall u\in\mathbf{U}(x).$$

Here one observes that the effective "discounting factor" is given by

$$\beta=\frac{T-1}{T}\in(0,1). \tag{14}$$

Then for any vectors $V_1, V_2:\mathbf{X}\to\mathbb{R}$,

$$\begin{aligned}
|\mathbf{T}[V_1](x)-\mathbf{T}[V_2](x)|\leq&\max_{u\in\mathbf{U}}\left|\Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u)+\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V_1(x')\right)-\Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u)+\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V_2(x')\right)\right|\\
\leq&\max_{u\in\mathbf{U}}\left|\max\left\{\mathcal{B}(x),\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V_1(x')-\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)V_2(x')\right\}\right|\\
\leq&\max_{u\in\mathbf{U}}\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)|V_1(x')-V_2(x')|\\
\leq&\max_{x'\in\mathbf{X}'}\frac{|V_1(x')-V_2(x')|}{\xi(x')}\max_{u\in\mathbf{U}}\sum_{x'\in\mathbf{X}'}\xi(x')\mathbb{P}(x'|x,u)\\
\leq&\max_{x'\in\mathbf{X}'}\frac{|V_1(x')-V_2(x')|}{\xi(x')}\beta\xi(x).
\end{aligned}$$

[3] $\|f\|_\xi=\max_{x\in\mathbf{X}}|f(x)|/\xi(x)$

This further implies that the following contraction property holds: $\|\mathbf{T}[V_1] - \mathbf{T}[V_2]\|_\xi \leq \beta \|V_1 - V_2\|_\xi$.

Similarly, the Bellman operator $\mathbf{T}_R$ also has the following properties.

Lemma 7 *The Bellman operator* $\mathbf{T}_R[V]$ *is monotonic, translational invariant and it is a contraction mapping with respect to the* $\|\cdot\|_\xi$ *norm.*

The proof of this lemma is identical to the proof of Lemma 6 and is omitted for the sake of brevity.

A.3 Proof of Theorem 2

The first part of the proof is to show by induction that for $x \in \mathbf{X}$,

$$V_N(x) := \mathbf{T}^N[V_0](x) = \min_\mu \Pi_{\mathcal{B}(x)} \left(\mathbb{E}\left[\sum_{t=0}^{N-1} \mathbf{D}(x_t, u_t) + V_0(x_N) \mid x, \mu \right] \right). \tag{15}$$

For $N = 1$, the definition of Bellman operator $\mathbf{T}$ implies that

$$V_1(x) = \mathbf{T}[V_0](x) = \min_{u \in \mathbf{U}(x)} \Pi_{\mathcal{B}(x)} (\mathbf{D}(x,u) + \mathbb{E}\left[V_0(x') \mid x,\ u\right]).$$

By the induction hypothesis, assume (15) holds at $N = k$. For $N = k+1$,

$$\begin{aligned}
&V_{k+1}(x) := \mathbf{T}^{k+1}[V_0](x) = \mathbf{T}[V_k](x) \\
&= \min_{u \in \mathbf{U}(x)} \Pi_{\mathcal{B}(x)} \left(\mathbf{D}(x,u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x,u) \left[\Pi_{\mathcal{B}(x')} \left(\min_\mu \mathbb{E}\left[\sum_{t=0}^{k-1} \mathbf{D}(x_t,u_t) + V_0(x_k) \mid x', \mu \right] \right) \right] \right) \\
&= \min_{u \in \mathbf{U}(x)} \max \left\{ \mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x,u) \left[\max \left\{ -\infty, \min_\mu \mathbb{E}\left[\sum_{t=0}^{k-1} \mathbf{D}(x_t,u_t) + V_0(x_k) \mid x', \mu \right] \right\} \right] \right\} \\
&= \min_{u \in \mathbf{U}(x)} \max \left\{ \mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x,u) \left[\min_\mu \mathbb{E}\left[\sum_{t=0}^{k-1} \mathbf{D}(x_t,u_t) + V_0(x_k) \mid x', \mu \right] \right] \right\} \\
&= \min_{u \in \mathbf{U}(x)} \Pi_{\mathcal{B}(x)} \left(\mathbf{D}(x,u) + \sum_{x' \in \mathbf{X}'} \mathbb{P}(x'|x,u) \left[\min_\mu \mathbb{E}\left[\sum_{t=1}^{k} \mathbf{D}(x_t,u_t) + V_0(x_{k+1}) \mid x', \mu \right] \right] \right) \\
&= \min_\mu \Pi_{\mathcal{B}(x)} \left(\mathbb{E}\left[\sum_{t=0}^{k} \mathbf{D}(x_t,u_t) + V_0(x_{k+1}) \mid x, \mu \right] \right).
\end{aligned}$$

Thus, the equality in (15) is proved by induction.

The second part of the proof is to show that $V^*(x_0) := \lim_{N\to\infty} V_N(x_0)$ and equation (5) holds. Since $V_0(x)$ is bounded for any $x \in \mathbf{X}$, the first argument implies that

$$\begin{aligned}
V^*(x_0) &= \min_\mu \max \left\{ 0, \lim_{N\to\infty} \mathbb{E}\left[\sum_{t=0}^{N-1} \mathbf{D}(x_t,u_t) + V_0(x_N) \mid x_0, \mu \right] \right\} \\
&\geq \min_\mu \max \left\{ 0, \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{D}(x_t,u_t) \,|\, x_0, \mu \right] \right\} - \lim_{N\to\infty} \max_{x \in \mathbf{X}'} \mathbb{P}[x_N = x \mid x_0, \mu] \|V_0\|_\infty \\
&\geq \min_\mu \max \left\{ 0, \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{D}(x_t,u_t) \mid x_0, \mu \right] \right\} - \epsilon \|V_0\|_\infty.
\end{aligned}$$

The first inequality is due to 1) V_0 is bounded and 2) $\mathbf{D}(x_t, u_t) = 0$ when x_t is in the absorbing set $\mathcal{X}$. The second inequality follows from the fact that x_t enters the absorbing set $\mathcal{X}$ after T steps. By similar arguments, one can also show that

$$V^*(x_0) \leq \min_\mu \max \left\{ 0, \mathbb{E}\left[\sum_{t=0}^{\infty} \mathbf{D}(x_t,u_t) \mid x_0, \mu \right] \right\} + \epsilon \|V_0\|_\infty.$$

Therefore, by taking $\epsilon \to 0$, the proof is completed.

The third part of the proof is to show the uniqueness of fixed point solution. Starting at $V_0 : \mathbf{X} \to \mathbb{R}$ one obtains from iteration $V_{k+1}(x) = \mathbf{T}[V_k](x)$ that

$$V_{k+1}(x) = \min_{u \in \mathbf{U}(x)} \max \left\{ \mathcal{B}(x), D(x,u) + \mathbb{E}\left[V_k(x') \mid x, u\right] \right\}.$$

By taking the limit, and noting that $V^*(x) = \lim_{k\to\infty} V_{k+1}(x) = \mathbf{T}[\lim_{k\to\infty} V_k](x) = \mathbf{T}[V^*](x)$, which implies V is a fixed point of the Bellman equation. Furthermore, the fixed point is unique because if there exists a different fixed point $\widetilde{V}$, then $\mathbf{T}^k[\widetilde{V}](x) = \widetilde{V}(x)$ for any $k \geq 0$. As $k \to \infty$, one obtains $\widetilde{V}(x) = V^*(x)$ which yields a contradiction.

A.4 Proof of Theorem 4

The convergence proof of two phase $Q-$learning is split into the following two steps.

Step 1 (Convergence of $Q-$update) We first show the convergence of $Q-$update (feasible set update) in two phase $Q-$learning. Recall that the state-action Bellman operator $\mathbf{F}$ is given as follows:

$$\mathbf{F}[Q](x,u) = \max\left\{\mathcal{B}(x), \mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q(x',u')\right\}.$$

Therefore, the $Q-$update can be re-written as

$$Q_{k+1}(x,u) = (1-\zeta_{2,k}(x,u))Q_k(x,u) + \zeta_{2,k}(x,u)\left(\Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_k(x',u')\right) + N_k(x,u)\right),$$

where the noise term is given by

$$N_k(x,u) = \Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u) + \frac{1}{N}\sum_{m=1}^{N} \min_{u'^{,m}\in\mathbf{U}(x'^{,m})} Q_k(x'^{,m},u'^{,m})\right) - \Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_k(x',u')\right), \tag{16}$$

for which $N_k(x,u)\to 0$ as $k\to\infty$ and for any $k\in\mathbb{N}$,

$$N_k^2(x,u) \le \left|\frac{1}{N}\sum_{m=1}^{N} \min_{u'^{,m}\in\mathbf{U}(x'^{,m})} Q_k(x'^{,m},u'^{,m}) - \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_k(x',u')\right|^2 \le 2\max_{x,u} Q_k^2(x,u).$$

Then the assumptions in Proposition 4.5 in [3] on the noise term $N_k(x,u)$ are verified. Furthermore, following the same analysis from Proposition 6 that $\mathbf{T}$ is a contraction operator with respect to the ξ norm, for any two state-action value functions $Q_1(x,u)$ and $Q_2(x,u)$, we have that

$$\begin{aligned}
&\left|\Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_1(x',u')\right) - \Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_2(x',u')\right)\right| \\
&\le \left|\sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_1(x',u') - \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}(x')} Q_2(x',u')\right| \\
&\le \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \max_{u'\in\mathbf{U}(x')} \left|Q_1(x',u') - Q_2(x',u')\right| \\
&\le \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u)\xi(x') \max_{x'\in\mathcal{X}} \max_{u'\in\mathbf{U}(x')} \frac{|Q_1(x',u') - Q_2(x',u')|}{\xi(x')} \le \beta\xi(x)\left\|Q_1 - Q_2\right\|_\xi.
\end{aligned} \tag{17}$$

Here $\|Q\|_\xi = \max_{x'\in\mathbf{X}} \max_{u'\in\mathbf{U}(x')} |Q(x',u')|/\xi(x')$ and $\beta\in(0,1)$ is given by (14) and ξ is given by (13). The first inequality is due to the fact that projection operator $\Pi_{\mathcal{B}(x)}$ is non-expansive. The second inequality follows from triangular inequality and

$$\sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u)\left|\min_{u'\in\mathbf{U}(x')} Q_1(x',u') - \min_{u'\in\mathbf{U}(x')} Q_2(x',u')\right| \le \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \max_{u'\in\mathbf{U}(x')} \left|Q_1(x',u') - Q_2(x',u')\right|.$$

The third inequality holds, due to the fact $\sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u)\xi(x') \le \beta\xi(x)$ for $\beta\in(0,1)$. Therefore the above expression implies that $\|\mathbf{F}[Q_1] - \mathbf{F}[Q_2]\|_\xi \le \beta\|Q_1 - Q_2\|_\xi$ for some $\beta\in(0,1)$, i.e., $\mathbf{F}$ is a contraction mapping with respect to the ξ norm.

By combining these arguments, all assumptions in Proposition 4.5 in [3] are justified. This in turns implies the convergence of $\{Q_k(x,u)\}_{k\in\mathbb{N}}$ to $Q^*(x,u)$ component-wise, where Q^* is the unique fixed point solution of $\mathbf{F}[Q](x,u) = Q(x,u)$.

Step 2 (Convergence of $H-$update) Now we show the convergence of $H-$update (objective function update) in two phase $Q-$learning. Since Q converges at a faster timescale than H, the $H-$update can be rewritten using the converged quantity, i.e., Q^*, as follows:

$$H_{k+1}(x,u) = H_k(x,u) + \zeta_{1,k}(x,u)\cdot\left(\mathbf{R}(x,u) + \frac{1}{N}\sum_{m=1}^{N} \min_{u'^{,m}\in\mathbf{U}_{\text{feas}}(Q^*,x'^{,m})} H_k(x'^{,m},u'^{,m}) - H_k(x,u)\right)$$

Recall that the state-action Bellman operator $\mathbf{F}_R$ is given as follows:

$$\mathbf{F}_R[H](x,u) = \mathbf{R}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}_{\text{feas}}(Q^*,x')} H\left(x',u'\right).$$

Therefore, the $H-$update can be re-written as the following form:

$$\begin{aligned}H_{k+1}(x,u) =&(1-\zeta_{1,k}(x,u))H_k(x,u)\\ &+\zeta_{1,k}(x,u)\left(\mathbf{R}(x,u)+\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)\min_{u'\in\mathbf{U}_{\text{feas}}(Q^*,x')}H_k\left(x',u'\right)+\mathcal{N}_k(x,u)\right),\end{aligned}$$

where the noise term is given by

$$\mathcal{N}_k(x,u)=\frac{1}{N}\sum_{m=1}^{N}\min_{u'^{,m}\in\mathbf{U}_{\text{feas}}(Q^*,x'^{,m})}H_k(x'^{,m},u'^{,m})-\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)\min_{u'\in\mathbf{U}_{\text{feas}}(Q^*,x')}H_k\left(x',u'\right), \tag{18}$$

such that $\mathbb{E}[\mathcal{N}_k(x,u)\mid\mathcal{F}_k]=0$ and for any $k\in\mathbb{N}$,

$$\begin{aligned}\mathcal{N}_k^2(x,u)\leq&\left|\frac{1}{N}\sum_{m=1}^{N}\min_{u'^{,m}\in\mathbf{U}_{\text{feas}}(Q^*,x'^{,m})}H_k(x'^{,m},u'^{,m})-\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)\min_{u'\in\mathbf{U}_{\text{feas}}(Q^*,x')}H_k\left(x',u'\right)\right|^2\\ \leq&2\max_{x,u}Q_k^2(x,u).\end{aligned}$$

Then the assumptions in Proposition 4.4 in [3] on the noise term $\mathcal{N}_k(x,u)$ are verified. Following the analogous arguments in (17), we can also show that $\|\mathbf{F}_R[H_1]-\mathbf{F}_R[H_2]\|_\xi\leq\beta\|H_1-H_2\|_\xi$ where $\beta\in(0,1)$ is given by (14) and ξ is given by (13), i.e., $\mathbf{F}_R$ is a contraction mapping with respect to the ξ norm. By combining these arguments, all assumptions in Proposition 4.4 in [3] are justified. This in turns implies the convergence of $\{H_k(x,u)\}_{k\in\mathbb{N}}$ to $H^*(x,u)$ component-wise, where Q^* is the unique fixed point solution of $\mathbf{F}_R[H](x,u)=H(x,u)$.

A.5 Proof of Theorem 5

The convergence proof of asynchronous two phase $Q-$learning is split into the following two steps.

Step 1 (Convergence of $Q-$update) Similar to the proof of Theorem 4, the $Q-$update in asynchronous two phase $Q-$learning can be written as:

$$Q_{k+1}(x,u)=(1-\zeta_{2,k}(x,u))Q_k(x,u)+\zeta_{2,k}(x,u)(\Theta_k(x,u)+\Psi_k(x,u)),$$

where

$$\Theta_k(x,u)=\begin{cases}\Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u)+\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)\min_{u'\in\mathbf{U}(x')}Q_k(x',u')\right) & \text{if } (x,u)=(x_k,u_k)\\ Q_k(x,u) & \text{otherwise}\end{cases}$$

and the noise term is given by

$$\Psi_k(x,u)=\begin{cases}N_k(x,u) & \text{if } (x,u)=(x_k,u_k)\\ 0 & \text{otherwise}\end{cases}$$

with N_k defined in (16). Since $N_k(x,u)\to0$ as $k\to\infty$, it can also be seen that $\Psi_k(x,u)\to0$ as $k\to\infty$. Furthermore, for any $k\in\mathbb{N}$, we also have that $\Psi_k^2(x,u)\leq N_k^2(x,u)\leq2\max_{x,u}Q_k^2(x,u)$. Then the assumptions in Proposition 4.5 in [3] on the noise term $N_k(x,u)$ are verified. Now we define the asynchronous Bellman operator

$$\widetilde{\mathbf{F}}[Q](x,u)=\begin{cases}\Pi_{\mathcal{B}(x)}\left(\mathbf{D}(x,u)+\sum_{x'\in\mathbf{X}'}\mathbb{P}(x'|x,u)\min_{u'\in\mathbf{U}(x')}Q(x',u')\right) & \text{if } (x,u)=(x_k,u_k)\\ Q(x,u) & \text{otherwise}\end{cases}.$$

It can easily checked that the fixed point solution of $\mathbf{F}[Q](x,u)=Q(x,u)$, i.e., Q^*, is also a fixed point solution of $\widetilde{\mathbf{F}}[Q](x,u)=Q(x,u)$. Next we want to show that $\widetilde{\mathbf{F}}[Q]$ is a contraction operator with respect to ξ. Let $\{\ell_k\}$ be a strictly increasing sequence ($\ell_k<\ell_{k+1}$ for all k) such that $\ell_0=0$, and every state-action pair (x,u) in $\mathbf{X}\times\mathbf{U}$ is being updated at least once during this time period. Since every state action pair is visited infinitely often, Borel-Cantelli lemma [26] implies that for each finite k, both ℓ_k and ℓ_{k+1} are finite. For any $\ell\in[\ell_k,\ell_{k+1}]$, the result in (17) implies the following expression:

$$\begin{aligned}&|\widetilde{\mathbf{F}}^{\ell+1}[Q](x,u)-Q^*(x,u)|\leq\beta\xi(x)\left\|\widetilde{\mathbf{F}}^{\ell}[Q]-Q^*\right\|_\xi && \text{if } (x,u)=(x_k,u_k)\\ &|\widetilde{\mathbf{F}}^{\ell+1}[Q](x,u)-Q^*(x,u)|=|\widetilde{\mathbf{F}}^{\ell}[Q](x,u)-Q^*(x,u)| && \text{otherwise}\end{aligned}$$

From this result, one can first conclude that $\widetilde{\mathbf{F}}[Q]$ is a non-expansive operator, i.e.,

$$|\widetilde{\mathbf{F}}^{\ell+1}[Q](x,u)-Q^*(x,u)|\leq\xi(x)\left\|\widetilde{\mathbf{F}}^{\ell}[Q]-Q^*\right\|_\xi.$$

Let $l(x,u)$ be the last index strictly between ℓ_k and ℓ_{k+1} where the state-action pair (x,u) is updated. There exists $\beta \in (0,1)$ such that

$$|\widetilde{\mathbf{F}}^{\ell_{k+1}}[Q](x,u) - Q^*(x,u)| \leq \beta\xi(x) \left\|\widetilde{\mathbf{F}}^{l(x,u)}[Q] - Q^*\right\|_\xi$$

From the definition of ℓ_{k+1}, it is obvious that $\ell_k < \max_{x,u} l(x,u) < \ell_{k+1}$. The non-expansive property of $\widetilde{\mathbf{F}}$ also implies that $\left\|\widetilde{\mathbf{F}}^{l(x,u)}[Q] - Q^*\right\|_\xi \leq \left\|\widetilde{\mathbf{F}}^{\ell_k}[Q] - Q^*\right\|_\xi$. Therefore we have that

$$|\widetilde{\mathbf{F}}^{\ell_{k+1}}[Q](x,u) - Q^*(x,u)| \leq \beta\xi(x) \left\|\widetilde{\mathbf{F}}^{\ell_k}[Q] - Q^*\right\|_\xi .$$

Combining these arguments implies that $\|\widetilde{\mathbf{F}}^{\ell_{k+1}}[Q] - Q^*\|_\xi \leq \beta \left\|\widetilde{\mathbf{F}}^{\ell_k}[Q] - Q^*\right\|_\xi$. Thus for $\delta_k = \ell_{k+1} - \ell_k > 1$ and $Q_k(x,u) = \widetilde{\mathbf{F}}^{\ell_k}[Q](x,u)$, the following contraction property holds:

$$\|\widetilde{\mathbf{F}}^{\delta_k}[Q_k] - Q^*\|_\xi \leq \beta \|Q_k - Q^*\|_\xi , \tag{19}$$

where the following fixed point equation holds: $\widetilde{\mathbf{F}}^{\delta_k}[Q^*](x,u) = Q^*(x,u)$. Then by Proposition 4.5 in [3], the sequence $\{Q_k(x,u)\}_{k\in\mathbb{N}}$ converges to $Q^*(x,u)$ component-wise, where Q^* is the unique fixed point solution of both $\mathbf{F}[Q](x,u) = Q(x,u)$ and $\widetilde{\mathbf{F}}[Q](x,u) = Q(x,u)$.

Step 2 (Convergence of $H-$update) Since Q converges at a faster timescale than H, the $H-$update in asynchronous two phase $Q-$learning can be rewritten using the converged quantity, i.e., Q^*, as follows:

$$H_{k+1}(x,u) = (1 - \zeta_{1,k}(x,u))H_k(x,u) + \zeta_{1,k}(x,u)(\Lambda_k(x,u) + \Phi_k(x,u)),$$

where

$$\Lambda_k(x,u) = \begin{cases} \mathbf{R}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}_{\text{feas}}(Q^*,x')} H_k(x',u') & \text{if } (x,u) = (x_k,u_k) \\ H_k(x,u) & \text{otherwise} \end{cases}$$

and the noise term is given by

$$\Phi_k(x,u) = \begin{cases} \mathcal{N}_k(x,u) & \text{if } (x,u) = (x_k,u_k) \\ 0 & \text{otherwise} \end{cases}$$

with $\mathcal{N}_k$ defined in (18). Since $\mathbb{E}[\mathcal{N}_k(x,u) \mid \mathcal{F}_k] = 0$, we have that $\mathbb{E}[\Phi_k(x,u) \mid \mathcal{F}_k] = 0$, i.e., $\Phi_k(x,u)$ is a Martingale difference. Furthermore we have that $\Phi_k^2(x,u) \leq \mathcal{N}_k^2(x,u) \leq 2\max_{x,u} Q_k^2(x,u)$ for $k \in \mathbb{N}$. The above arguments verify the assumptions in Proposition 4.4 in [3] on the noise term $\Phi_k(x,u)$. Now define the asynchronous Bellman operator

$$\widetilde{\mathbf{F}}_R[H](x,u) = \begin{cases} \mathbf{R}(x,u) + \sum_{x'\in\mathbf{X}'} \mathbb{P}(x'|x,u) \min_{u'\in\mathbf{U}_{\text{feas}}(Q^*,x')} H(x',u') & \text{if } (x,u) = (x_k,u_k) \\ H(x,u) & \text{otherwise} \end{cases} .$$

It can easily checked that the fixed point solution of $\mathbf{F}_R[H](x,u) = H(x,u)$, i.e., H^*, is also a fixed point solution of $\widetilde{\mathbf{F}}_R[H](x,u) = H(x,u)$. Then following analogous arguments from step 1 (in particular expression (19)), for $\delta_k = \ell_{k+1} - \ell_k > 1$ and $H_k(x,u) = \widetilde{\mathbf{F}}^{\ell_k}[H](x,u)$, one shows that $\|\widetilde{\mathbf{F}}^{\delta_k}[H_k] - H^*\|_\xi \leq \beta \|H_k - H^*\|_\xi$ for some $\beta \in (0,1)$, which further implies the following fixed point equation holds: $\widetilde{\mathbf{F}}^{\delta_k}[H^*](x,u) = H^*(x,u)$. Thus by Proposition 4.4 in [3], the sequence $\{H_k(x,u)\}_{k\in\mathbb{N}}$ converges to $H^*(x,u)$ component-wise, where Q^* is the unique fixed point solution of both $\widetilde{\mathbf{F}}_R[H](x,u) = H(x,u)$ and $\mathbf{F}_R[H](x,u) = H(x,u)$.